

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Dynamic Allocation (DA)</b>	<p>140. Google, mirroring the duties of financial brokers to their clients, promised publishers that its publisher ad server would act in their best interests. Google told publishers, for instance, that Dynamic Allocation maximized their inventory yield; it ‘maximized revenue,’ Google advertised about its publisher ad server. Google also told publishers that, with Dynamic Allocation, publishers have a ‘risk-free way to get the highest real-time revenues for all their non-guaranteed impressions.</p> <p>141. In fact, Google concealed the nature of its conduct and knew that Dynamic Allocation did not in fact maximize publishers’ yield. Google internally discussed how publishers could make more money selling their inventory if exchanges really competed. Internal Google documents reveal Google’s knowledge of its own misrepresentations, stating that ‘the optimal publisher set up includes multiple exchanges in order to capture the largest demand pool and increase RPMs [revenue per impression] through [exchange] competition.’ In fact, according to one Google study, competition between exchanges increased publishers’ clearing prices by an average of 40 percent. In other words, Dynamic Allocation had permitted Google’s exchange to clear publishers’ inventory for depressed prices. One industry publication put it succinctly, “[t]he lack of competition was costing pub[s] cold hard cash.”</p>	<p>278. Unsurprisingly, Google concealed the nature of its conduct and did not tell publishers the truth about Dynamic allocation. <u>Google induced publishers to use DFP</u> (including DFP’s Dynamic allocation ‘feature’) by representing that it maximized publishers’ inventory yield. Google advertised to publishers, for instance, that DFP/Dynamic Allocation ‘maximizes revenue.’ Google also told publishers that, with Dynamic Allocation, publishers have a ‘risk-free way to get the highest real-time revenues for all their non-guaranteed impressions.”</p> <p>279. Google, however, knew that Dynamic Allocation did not maximize publishers’ yield. <u>Internal Google documents reveal Google’s knowledge of its own misrepresentations, stating that the optimal publisher set up in display advertising includes “real-time bidding across exchanges,” which is “at scale, at the best possible price, with zero waste.”</u></p>
<b>Enhanced Dynamic Allocation (EDA)</b>	<p>182. Similar to Google’s strategy with Dynamic Allocation, Google invited publishers to enable EDA under a false pretense. Wearing their publisher ad server hat, Google falsely told publishers that EDA ‘maximizes yield.’ EDA did not, however, maximize publishers’ yield. Internally, Google understood that the EDA program was a scheme to let</p>	<p>291. Google automatically turned EDA on for publishers then coaxed publishers into leaving EDA turned on under a false pretense. Wearing its publisher ad server hat, Google falsely told publishers that EDA ‘maximizes yield.’ <u>Publishers relied upon Google’s misrepresentations to enable EDA, thinking it would maximize yield.</u></p>

<sup>1</sup> Underlining indicates key language that was added to the TAC.

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Enhanced Dynamic Allocation (EDA)</b> <i>(continued)</i>	<p>Google's exchange simply 'cherry-pick [publishers'] higher-revenue impressions.' In fact, cherry-picking the best impressions under EDA helped Google make an additional \$150 million per year.</p> <p>183. To make matters worse, Google's practice of scrambling user IDs (discussed above in paragraphs 142-147) concealed the true nature of Google's conduct. Publishers could not easily know that, with EDA, Google was cherry-picking impressions. By scrambling the IDs differently for publishers and advertisers, publishers could not easily work with advertisers to confirm that advertisers were receiving the valuable impressions (e.g., ads shown to users with high net worth) as opposed to the low value ones (e.g., ads shown to a 10-year-old child with no purchasing power).</p>	<p>292. <u>Google knew EDA did not and would not actually maximize publishers' yield.</u> Internally, Google understood that EDA was a scheme to let its own AdX exchange simply 'cherry-pick [publishers] higher-revenue impressions,' earning Google's exchange an additional \$150 million per year. Moreover, Google concealed the true nature of its conduct by hashing publishers' server IDs.</p>
<b>Project Bernanke</b>	<p>151. In 2013, Google's gTrade team designed Project Bernanke, yet another program to exclude competition. Named after the former Federal Reserve Chairman, Project Bernanke uses privileged access to detailed information regarding what advertisers historically bid to help advertisers using Google Ads beat the advertisers bidding through competitors' ad buying tools...The Bernanke program is designed so that it is not transparent to publishers.</p>	<p>298. In 2013, Google's gTrade team devised and launched a <u>secret</u> program, which they codenamed 'Project Bernanke' (after the quantitative easing policy of the former Federal Reserve Chairman). <u>Project Bernanke deceives publishers and advertisers and excludes competition in the exchange market and the buying tools for small advertisers market.</u></p> <p>299. By way of background, "first-price" and "second-price" auctions are common types of auctions used in various industries and contexts. Generally speaking, in a first-price auction, the buyer pays the amount of their own winning bid; and as the name implies, the buyer in a second-price auction pays the amount of the second-highest bid (sometimes with a negligible additional amount, e.g., one penny). A "third-price" auction, therefore, is one in which the buyer pays the amount of the <i>third</i>-highest bid. As addressed below, Google's secret Bernanke program surreptitiously switched Google's AdX exchange from a second-price auction to a third-price auction</p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Project Bernanke</b> <i>(continued)</i>		<p>on billions of impressions per month.</p> <p>300. <u>Between 2010 and September 2019, Google led publishers and advertisers to believe that AdX was a second-price auction. For example, shortly after launching its AdX exchange in 2009, Google executive Scott Spencer promoted the new product by discussing AdX's mechanics in an interview published on the popular industry website AdExchanger. Spencer explained: "AdX is a second price auction with minimum CPMs set by the publisher. This is the most efficient auction model, resulting in the most stable, long-term equilibrium price." In their 2014 paper "Yield Optimization of Display Advertising with AdExchange" (published in the <i>American Economic Review</i>), Google senior researchers Jon Feldman, Vahab Mirrokni, and S. Muthukrishnan similarly promoted AdX: "With multiple bidders, AdX runs a sealed bid second-price auction." No doubt, publishers and advertisers were led by Google to believe that when AdX ran an auction, the highest bidder would win and pay the amount of the second-highest bid.</u></p>
<b>Dynamic Revenue Share (DRS)</b>	<p>150. Google's gTrade team launched another program called Dynamic Revenue Share (DRS) that leverages exclusive access to publishers' ad server user IDs to exclude exchange competition in a second way. Google automatically opted publishers into the DRS program under the misrepresentation that it would make publishers more money . . . .</p>	<p>324. <u>Google internally acknowledged that DRS made its auction untruthful: "One known issue with the current DRS is that it makes the auction untruthful as we determine the AdX revshare after seeing buyers' bids and use winner's bid to price itself (first-pricing)...."</u></p> <p>326. <u>Not surprisingly, Google concealed DRS from both publishers and advertisers. Google started opting publishers into DRS starting in 2014 without disclosing anything about the program to publishers or advertisers. By the fall of 2015, Google had opted all publishers into DRS, still without disclosing the program.</u></p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Dynamic Revenue Share (DRS)</b> <i>(continued)</i>		<p>327. <u>In the summer of 2016, without referring to the program’s real name, Google told publishers it was launching a “revenue share-based optimization” that increased a publisher’s yield. Google was referring to DRS, which plainly did not increase publisher yield.</u></p> <p>328. <u>Google continued to mislead publishers and advertisers about the program and withheld critical information that the parties could have used to make an informed decision about the program. For instance, Google did not disclose that Google had been operating the program since 2014 or that DRS set floors on a post-hoc basis after peeking at received bids. And this was the critical point. Peeking ahead at other exchanges’ net bids, then altering AdX’s margin after peeking at its own received bids, permitted AdX to win when it should have been excluded according to publishers’ wishes.</u></p>
<b>Reserve Price Optimization (RPO)</b>	<p>149. Google’s New York-based quantitative team “gTrade” designed one such program called Reserve Price Optimization (“RPO”). Google’s RPO program uses exclusive access to publishers’ user IDs to dynamically adjust the price floors in Google’s exchange on a per-buyer basis depending on what Google knows a particular buyer will actually pay . . . .</p>	<p>340. Google launched RPO in early 2015 and automatically opted publishers into the program without their knowledge or consent.</p> <p>341. <u>Around the same time, Google publicly and falsely denied plans to launch dynamic floors in its exchange. On March 5, 2015, Digiday ran a short story based on a leak about Google’s potential plans to launch dynamic price floors. The publication asked Google whether it planned to adjust price floors based on publishers’ use of Google’s DFP ad server. In response, spokeswoman Andrea Faville issued a statement: “That description doesn’t match anything in our current product suite or future roadmap.” Ms. Faville’s statement directly contradicted Google’s internal operations. Internally, Google planned to launch RPO weeks later for 50 percent of publishers by April 7 and for 90 percent by April 17.</u></p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Reserve Price Optimization (RPO)</b> <i>(continued)</i>		<p>342. <u>Instead, Google continued to mislead publishers by encouraging them to adjust Google exchange floors in their publisher ad server. DFP continued to let publishers pre-set floors for Google’s AdX exchange, buying tools, and advertisers, directly leading them to believe that they could control outcomes and optimize yield through floors.</u></p> <p>343. <u>Over a year later, on May 12, 2016, Google announced it was launching ‘optimized pricing.’ However, Google did not disclose that it had actually launched RPO over a year earlier, did not disclose that RPO relied on inside information, and misled publishers and advertisers as to how the program worked.</u></p> <p>344. <u>For instance, Google’s publisher-facing team told publishers that the dynamic floor program increased inventory yield. On the surface, RPO appeared to increase yield because AdX initially returned higher bids. However, because RPO relied on inside information, combining bid data from AdX with publishers’ ad server user IDs, it exacerbated problems of adverse selection in publishers’ inventory auctions. Markets rife with problems of adverse selection are inefficient, dissuade participants from entering, and result in lower output. As a result, RPO ultimately forecloses competition from exchanges and advertisers and reduces inventory yield. Publishers could not discover this harmful effect because Google failed to disclose RPO’s reliance on inside information.</u></p> <p>345. <u>In addition to misleading publishers, Google misled advertisers and misrepresented how the program worked. In its blog post disclosing RPO, Google claimed that it would “monitor [optimized pricing’s] performance to ensure advertisers continue[d] to get great ROI” and that it would “give programmatic buyers greater access to premium</u></p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Reserve Price Optimization (RPO)</b> <i>(continued)</i>		<p><u>inventory.” Google also approached select large, sophisticated buyers on a one-on-one basis representing that dynamic floors were good for them. Google kept a record of these conversations and advertiser responses. According to Google’s records, one advertiser pushed back, asking “How is this good for the buyer? Because I’ll tell you, it isn’t. It just raises the price.” Google responded misleadingly by saying that the program helps advertisers by increasing the amount of inventory available for purchase programmatically. Privately, employees acknowledged that RPO did not help advertisers at all.</u></p> <p><u>346. All the while, Google continued to lead publishers and advertisers to believe that AdX operated a second-price auction, inducing advertisers to submit a sealed bid reflecting their true value. Numerous industry articles covering Google’s conduct in the exchange market continued to report that Google operated a second-price auction. Internally, Google employees discussed public perception around AdX operating as a second-price auction. It was not until 2019 that Google publicly migrated to a first-price auction, discarding all pretense of running a second-price auction.</u></p> <p><u>347. Google’s internal documents reveal that Google was aware of the resulting deception and harm in the market. In an email between colleagues discussing RPO, a Google employee wrote: “Doesn’t that undermine the whole idea of second price auctions? I.e., the assurance that you can bid the maximum you’re willing to pay with no negative consequence. But if the publisher manufactures a floor based on your bid to get you to pay more than the second price, this principle gets violated. It’ll transform the system into a 1st price auction where the bidder has a strong incentive to bid LESS than he’s willing to pay. (Only just enough to win.) I don’t think that’s desirable for</u></p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Reserve Price Optimization (RPO)</b> <i>(continued)</i>		<u>either side in the long term.</u> Another employee wondered: “ <u>Is RPO not just basically pushing our second price auction - that is supposed to be fair - toward a first priced auction?</u> ”
<b>Exchange Bidding</b>	<p>200. Fourth, Google foreclosed exchange competition by secretly rigging the Exchange Bidding program to let Google win. Google designed Exchange Bidding to provide Google’s exchange a special “prioritization,” which Google kept secret. Google made it so its own AdX exchange won publishers’ inventory even over another exchange’s higher bid. In the following email, Google employee [redacted] explained how the Exchange Bidding program returned results that were “suboptimal for pubs yield”: a Google AdX bid of \$6 would win even though another exchange (“EB SSP”) submitted a higher \$8 bid.</p> <p>201. Internally, Google employees grappled with the fact that Google was falsely telling publishers that Google’s header bidding alternative enabled competition and improved yield, since in reality, Google created a program that advantaged itself at the expense of publishers. As one senior Google employee observed, Exchange Bidding’s deliberate design is to avoid price competition, which “generates substantial yields for publishers and serious risks of negative media coverage if exposed externally.”</p>	<p>379. Several years later, in 2019, Google publicly announced that exchanges in Exchange Bidding would no longer be able to trade ahead of header bidding exchanges. Google represented that it would run a “a [sic] fair and transparent market for everyone.” It said that “every offer from programmatic buyers will compete in the same unified auction, alongside inventory which is directly negotiated with advertisers. An advertising buyer’s bid will not be shared with another buyer before the auction or be able to set the price for another buyer.” But Google’s auction was neither fair nor transparent. Rather, Google continued to advantage its own exchange and other Exchange Bidding exchanges over header bidding exchanges.</p>
<b>Header Bidding</b>	<p>241. Internal communications between Google employees reveal how Google engaged in deception to undermine header bidding and foreclose competition in the exchange market. In one instance, the OpenX exchange noticed their auction transactions and revenue in header bidding plummet. When OpenX reached out to a publisher to diagnose the problem, the publisher explained that Google employees told the publisher to remove the OpenX exchange from header bidding to solve a</p>	<p>385. Internal communications between Google employees reveal that Google engaged in further deception to undermine header bidding and foreclose competition in the exchange market. <u>Indeed, Google had launched the “Header Bidding Observatory,” an internal division at Google formed to develop strategies to combat header bidding, which detected whether publishers had enabled header bidding and launched “proactive outreach” to “optimize these implementations.”</u> In one instance,</p>



## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Header Bidding</b> <i>(continued)</i>	“strain on its servers” and improve the publisher’s yield. However, a senior Google employee worried its misrepresentations would make it difficult “to convince [companies] to trust us.” Another employee conceded it gave Google a “bad look.” Google employees agreed that, in the future, they should find ways to convince publishers to act against their interest and remove competing exchanges in header bidding on their own.	the OpenX exchange noticed their auction transactions and revenue for a particular major New York publisher in header bidding plummet. When OpenX reached out to the publisher to diagnose the problem, the publisher explained to OpenX that Google had advised them to remove the OpenX exchange from header bidding to solve a “strain on its servers” and improve the publisher’s inventory yield. <u>Google used its position as publishers’ agent to deceive publishers to act against their own interests.</u> Internally at Google, a senior employee worried that Google’s misrepresentations would make it difficult “to convince [companies] to trust us.” Another employee conceded it gave Google a “bad look.” Google employees agreed that, in the future, they should find ways to convince publishers to act against their own interests and remove competing exchanges in header bidding on their own.
<b>Accelerated Mobile Pages (AMP)</b>	248. After crippling AMP’s compatibility with header bidding, Google went to market falsely telling publishers that adopting AMP would enhance page load times. But Google employees knew that AMP only improves the “median of performance” and AMP pages can actually load slower than other publisher speed optimization techniques. In other words, the ostensible benefits of faster load times for a Google-cached AMP version of a webpage were not true for publishers that designed their web pages for speed. Some publishers did not adopt AMP because they knew their pages actually loaded faster than AMP pages.	411. Publicly, Google presented the true objectives of AMP as decreasing page load time and reducing latency. <u>Internally, however, Google identified that AMP was an avenue to combat the adoption of header bidding.</u>
<b>Jedi Blue / Unified First Price Auction</b>	233. As one would expect with a market allocation agreement, Google and Facebook do not disclose their secret match rate, bid rate, or win rate agreements to other auction participants. Rather, Google publicly misrepresents that all bidders in publishers’ auctions compete on equal footing. “All participants in the unified auction, including Authorized	447. Given the choice, no rational developer would choose to have its auctions rigged by the market’s two largest buyers. So, Google and Facebook swore themselves to secrecy about the terms of their agreement and have not generally disclosed their secret match-rate, bid-rate, or win-rate agreements to either developers or other auction participants. <u>They have had plenty</u>



## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Jedi Blue / Unified First Price Auction</b> <i>(continued)</i>	<p>Buyers and third-party yield partners, compete equally for each impression on a net basis,” Google publicly markets on its website. This, of course, is patently false. It is false not only because of the special terms in the Facebook agreement, but also because Google used algorithms to systematically manipulate auction outcomes and repeatedly traded on inside information to win auctions.</p>	<p><u>of opportunity to do so: implementing the Jedi Blue agreement requires both Google and Facebook to update and re-execute their respective agreements with the app developers whose advertising inventory they hope to purchase. When encouraging developers to update to the latest version of their respective bidding agreements, Google and Facebook obfuscate the true motive for the contract changes, which say nothing about the auction manipulation terms of Jedi Blue. Google uses the promise of competitive bidding between Google and Facebook’s ad networks as an inducement for developers to sign new contracts, when in fact Jedi Blue secretly limits the terms of the competition between those two bidders.</u></p> <p>448. <u>Google also advertises the ability to accept Facebook bids as a feature of its mediation tool and has used the promise of competitive bidding by rival ad networks as a lure to further increase its share of the in-app mediation market. But this is a classic bait and switch. Rather than robust competition between Google and Facebook’s ad networks, an app developer adopting Google’s mediation tool gets a form of sham competition between those supposed rivals, one that is constrained by secret terms in Jedi Blue. And once a developer has adopted Google’s mediation tool, the costs of switching to a different tool are substantial. Doing so would require rewriting source code, implementing a new tool, and testing that tool’s compatibility with various ad networks. For the very large number of developers who have adopted Google’s mediation tool—representing more than half of all applications featuring any advertising—Google now has an ability to impose anticompetitive and collusive terms that developers never would have accepted in advance.</u></p>
<b>Encrypting User IDs</b>	<p>156. Google’s publicly stated reason for its publisher ad server cutting off publishers’ ability to share their ad server user IDs</p>	<p>482. Google’s publicly stated reason for cutting off publishers’ ability to share ad server user IDs is the purported protection of</p>

## APPENDIX A

Topic	Second Amended Complaint	Third Amended Complaint <sup>1</sup>
<b>Encrypting User IDs</b> <i>(continued)</i>	with non-Google exchanges is the purported protection of users' privacy. But Google does not actually care about users' privacy. Rather, Google wants to prevent companies from creating deeper and more comprehensive user profiles by combining different sets of user data. However, Google's ad server shares those very user IDs with Google's exchange and buying tools. Google then does what it wants to prevent others from doing: it combines the data sets to create more comprehensive user profiles and deliver more targeted advertising.	users' privacy. <u>Specifically, Google claims that encrypting IDs reduces companies' ability to combine data sets to create more intrusive individual user data profiles. However, Google's ad server has no qualms with sharing user IDs amongst Google's own buying tools and with Facebook through its Jedi Blue agreement. This allows Google to do precisely what it seeks to prohibit others from doing: combining user data sets in secret databases with codenames like HULK and MindReader to deliver more targeted advertising.</u> Therefore, contrary to Google's purported concern for user privacy, Google prevents publishers from providing their consumers with similar privacy benefits from Google's own network, exchange, and buying tools.